

Response Latency and Candidate Availability in High-Volume Hourly Hiring

A structured review of the evidence, and a framework for measuring the interval no published study has measured

Patrick Underwood

Kordis Research, Toronto, Ontario

15 July 2026

This paper is issued in the Kordis Research working paper series. It has not been peer reviewed and should be cited as a working paper. Kordis develops operations software for temporary staffing agencies and therefore has a commercial interest in several of the questions examined here; the disclosure at the end of the paper states that interest in full, together with the source verification protocol. No external funding was received and no client data was used. Correspondence: research at kordisai.com.

Abstract

High-volume hourly hiring is managed almost everywhere as an inventory problem in which the inventory does not deteriorate. Recruiting effort is directed at raising the number of applications, and the resulting pool is worked at whatever rate office capacity allows. We ask whether the probability that an applicant converts to a completed first shift declines convexly in the elapsed time between application and first substantive employer contact, and whether the operative scarce resource is therefore the depth of the applicant pool or the availability remaining in it. A structured review of the peer-reviewed, statistical-agency and practitioner literature returns a negative result: no peer-reviewed study measures this interval, no estimate of the functional form of the decline exists, and the applicant withdrawal literature does not model time. The universal practitioner prescription to respond quickly rests on assertion. We assemble what the adjacent literatures do establish, formalise candidate availability as a time-dependent quantity, distinguish three functional forms that current guidance conflates, and specify the event schema and identification strategy for the hazard study that would settle the question.

Keywords: recruitment; response latency; applicant withdrawal; high-volume hiring; light industrial staffing; temporary help services; selection validity; hazard models; Canada

JEL classification: J23, J63, M51, M55

1 Introduction

Between October 2024 and October 2025, the number of applicants per open manufacturing role in the United States rose from 41 to 47. Over the same twelve months, hiring in manufacturing fell by 7 per cent, openings rose by 14 per cent, and the time taken to fill a role lengthened from 40 days to 42 (iCIMS 2025). Supply went up. Conversion went down. The process got slower.¹

Set aside for a moment the question of what caused each of those three movements. Consider only whether they can be true at once under the model of hiring that almost every high-volume operation actually runs. That model treats the applicant pool as a stock. Marketing spend raises the stock. Recruiter capacity determines the rate at which the stock is worked. If the stock rises and the rate is unchanged, hires should rise, or at worst hold. They should not fall.




Figure 1. The high-volume hiring paradox in US manufacturing. Applicant supply per opening, the number of openings, and the time taken to fill them all rose over the year. Hires fell. Under a model in which the applicant pool is a static stock worked at a fixed rate, more supply cannot produce fewer hires. Under a model in which the pool decays while it waits, it can, because more supply at fixed processing capacity means an older pool at the moment of contact.

Source: Compiled from iCIMS, *November 2025 Workforce Report*, reporting October 2025 against October 2024. Vendor publication. Openings and hires are as reported. Applicants per opening (41 to 47) and days to fill (40 to 42) are reported as levels; the percentage changes shown are computed from those levels for comparability on a single axis.

The pattern resolves immediately under a different assumption. Suppose an application is not a unit of inventory but a claim with an expiry. Suppose that what an employer holds after someone applies is not a candidate but an option on a candidate, and that the option decays. Then raising the number of applications while holding processing capacity fixed does not raise hires. It raises the average age of the pool at the moment of contact, and if conversion falls faster in age than volume rises, hires fall. The observation that supply and conversion moved in opposite directions is what a model with time-dependent availability predicts.

This paper takes that possibility seriously enough to ask what evidence exists for it, and finds that the answer is: less than the confidence of industry practice would suggest. The specific research question is:

1. These figures are vendor-published and drawn from one applicant tracking platform's customer base. They are reported here because no government statistical agency in Canada or the United States publishes applicant volume per opening, and because the direction of the pattern, not its precise magnitude, is what the argument requires.

Research question. In high-volume hourly and light industrial hiring, does the probability that an applicant converts to a completed first shift decline convexly in the elapsed time between application submission and first substantive employer contact? If it does, is the operative scarce resource the depth of the applicant pool or the availability remaining in it?

Two features of that question are deliberate. The first is the outcome. The terminal event is a completed first shift, not a hire. In temporary staffing this is not a refinement, it is the definition of delivery: under Ontario law an agency worker is an assignment employee whose employment relationship with the agency persists between assignments, so a placement that produces no worked shift is a record, not a unit of output (Ontario, Ministry of Labour). The second is the interval. The clock starts when the applicant submits and stops at first substantive contact, because that is the interval an operator actually controls. It is also, as section 3 shows, the interval nobody has measured.

2 Evidence base and review protocol

This paper is a structured evidence review. It is not a systematic review in the PRISMA sense, and it is not a statistical meta-analysis: we do not pool effect sizes into a summary estimate, because the studies bearing on the question measure different constructs on different populations and a pooled estimate would be uninterpretable. Where several studies report an estimate of the same parameter we tabulate them side by side with the number of independent samples and the total participants, and we let the reader see the dispersion rather than average it away. The protocol is set out here so that the review can be audited and repeated.

2.1 Scope and inclusion

The scope is the recruitment and selection literature in industrial and organisational psychology and human resource management, the labour economics of search and matching, official Canadian and United States labour statistics, Ontario employment standards law as it governs assignment employees, and the practitioner and vendor literature that circulates operating advice to staffing firms. A source was included if it bore directly on one of the constructs in the research question, or if it established the absence of evidence on one of them. Recency was not an inclusion criterion: the foundational results in several of these literatures are decades old and remain the best available, and where a newer estimate supersedes an older one we report both and say which is current.

2.2 Search

The search was iterative rather than a single indexed query. It proceeded from the constructs in the research question outwards, through citation chains, through the reference lists of the meta-analyses and reviews retrieved, and through the publication lists of the statistical agencies and regulators with jurisdiction over the setting. Because the search was iterative, we do not report a records-screened count; a screening count implies a single reproducible

query, and reporting one here would overstate the formality of the method. This is the principal respect in which the review falls short of a systematic review, and it is stated rather than concealed. Searching concluded in July 2026.

2.3 Verification

Every source cited was retrieved and read at the address given in the reference list. No figure, effect size, quotation, statutory provision or interval appears anywhere in this paper unless it was read in a retrieved document. Where a publisher platform refused access, we substituted an institutional repository, an author-hosted manuscript, or a bibliographic index record, and where only an abstract could be reached we report the design without reporting quantities that the abstract does not carry. Sources we expected to exist and could not verify were omitted rather than cited from memory, and the substantive omissions are named in the limitations section. Two quotations from the same practitioner source are set as separate quoted fragments rather than joined, because the original punctuation between them could not be reproduced under the typographic rules of this series.

2.4 Classification

Sources are classified by publication type on retrieval, and the class travels with the source: it is stated in the note to any figure or table that draws on it, and the full assignment is tabulated in the appendix. The distinction that matters most here is between evidence produced by parties with a commercial interest in the answer and evidence produced by parties without one. Of the 35 sources in this paper, 19 are peer-reviewed or peer-reviewed-adjacent, 8 are from statistical agencies, government departments or regulators, and 8 are from industry bodies, research centres, or vendor and practitioner publications.

Table 1. Composition of the evidence base, by source class.

Source class	Sources	Share (%)
Peer-reviewed	17	49
Working paper	1	3
Doctoral dissertation	1	3
National statistical agency	5	14
Government department or agency	2	6
Regulator or statute	1	3
Industry body	1	3
Vendor or practitioner publication	6	17
Author's own research programme	1	3
Total	35	100

Notes: Classification is by publication type, assigned on retrieval and not revised afterwards. “Peer-reviewed” covers journal articles and peer-reviewed volume chapters; working papers, dissertations and open courseware are counted separately because they carry a different evidential weight. “Regulator or statute” covers a regulator’s own

published guidance and the text of a regulation. “Vendor or practitioner publication” covers material published by firms selling into the market under study, and by practitioner magazines; every figure taken from that class is identified as such at the point of use, in the figure or table note as well as in this appendix.

2.5 Negative findings

This review reports five findings of the form “this has not been measured.” Such a claim is weaker than a claim about the world and we mark the difference throughout: what we assert is that we searched a stated body of literature by a stated method and did not find the result, not that the result does not exist. Where a retrieved document was examined specifically to confirm that it does not contain a quantity, that is recorded as a verified negative and the document is cited for the absence.

3 What is known, and what is not

3.1 The prescription without the evidence

Advice to respond to applicants quickly is close to universal, and it is striking how thinly it is sourced. The Society for Industrial and Organizational Psychology’s own practitioner synthesis instructs employers to “make timely feedback a priority and track time-to-feedback for each selection hurdle,” and reports no effect size and no benchmark for latency anywhere in the document (SIOP 2012). The most widely circulated operational rule, that a response after two days is a response too late, appears in vendor guidance with no sample size, no methodology and no definition of the threshold (Workable 2014). The most cited statistic on candidate patience, that two thirds of jobseekers abandon after two weeks, is a self-reported intention with no published sample or method (CareerBuilder n.d.).

Table 2. What the practitioner literature prescribes about response speed, and what evidence stands behind each prescription.

Prescription as circulated	Stated basis	Strongest verifiable support
Track time to feedback at every selection hurdle	Professional body guidance	None given. The source document reports no effect size and no benchmark for latency (SIOP 2012).
Respond within two days or lose the candidate	Asserted as a rule of thumb	No sample, no methodology, no threshold definition are published (Workable 2014).
Candidates give up after about two weeks	Vendor candidate-behaviour survey	Self-reported intention rather than observed behaviour; no sample size or method published (CareerBuilder n.d.).
Automating the screening call cuts interview no-shows sharply	Vendor product claim	No sample, no methodology, and no reported relationship between days-to-interview and no-show probability (Workstream 2026).

Prescription as circulated	Stated basis	Strongest verifiable support
Delays signal a poorly run employer	Applicant reactions research	Supported in direction. Recruitment delays had “strongly negative effects” and were read as symbolic of the organisation (Rynes et al. 1990); process and recruiter factors predict attraction most strongly early in the funnel (Uggerslev et al. 2012).
Moving faster costs you candidate quality	Folk theory	Contradicted. Faster offers raised acceptance with no penalty to performance ratings or turnover across 3,012 candidates (Becker et al. 2010).

Notes: Classification follows the source’s own publication type. The point of the table is not that practitioner guidance is wrong. It is that the guidance is universal and the evidence base under it is close to empty, which is what makes the measurement gap in section 2.3 worth closing.

None of this means the advice is wrong. Section 4 gives reasons to think the direction is right. It means that a claim which is treated as settled in practice is unsettled in the record, and that the gap between the two is large enough to be worth measuring rather than asserting.

3.2 *What the peer-reviewed record contains*

The peer-reviewed literature on applicant reactions is substantial and mature. Its central finding, established across 86 independent studies and 48,750 participants, is that applicants who perceive a selection process as fair view the organisation more favourably and report stronger intentions to accept (Hausknecht et al. 2004). The framework that organises this work treats selection as a procedural justice problem (Gilliland 1993), and the measurement instruments built on it operationalise fairness in terms of procedure, explanation and treatment (Bauer et al. 2001).

What this literature does not code is latency. The meta-analysis catalogues procedure type and perceived validity; response time is not among its moderators.² The nearest thing to a timing result comes from an archival study of 3,012 candidates at a single large employer, which found that candidates of both student and experienced type were more likely to accept earlier offers, and that faster offers produced no penalty in performance ratings or turnover (Becker et al. 2010). That study measures the interview-to-offer interval in a professional hiring setting, not the application-to-contact interval in an hourly one, but its second finding matters independently: it removes the standard defence that a slower process is a more careful one.

Two further results bear on the shape of the problem rather than its magnitude. Recruitment delay was found to have “strongly negative effects” on job choice, and, importantly, to be read by candidates as symbolic of the organisation as a whole, with the effect concentrated among candidates who had greater search success and therefore more alternatives (Rynes et al. 1990). And a meta-analysis spanning 232 studies, 250 samples and 108,632

2. We looked for a timeliness moderator specifically, and did not find one. This is a claim about what the published moderator lists contain, and it should be read as such.

participants found that recruiter and process factors carry substantial incremental predictive power for applicant attraction early in the recruiting process and much less later, while job characteristics dominate at the final choice (Uggerslev et al. 2012). If process conduct matters most early, then the cost of latency should be front-loaded, which is the convexity this paper asks about. A recent experimental study reaches a compatible conclusion from the other direction, reporting across two studies of 523 candidates that timely communication mattered most during the initial selection stages (Barattucci et al. 2025).

3.3 Four gaps, stated precisely

We searched for each of the following and did not find it. These are claims about the state of publication, not about the state of the world.

1. *No peer-reviewed study measures application-to-first-contact latency as a predictor of hire.* The closest is Becker et al. (2010), which measures a later interval in a different labour market.
2. *No published estimate of the decay's functional form exists.* Convexity against employer latency has not been tested in a hiring setting. The one hiring study that establishes convex decay with a flat tail runs its clock on the applicant's unemployment duration (Kroft et al. 2013); the one study that establishes convex decay against response latency measures a qualified sales lead rather than a hire (Oldroyd et al. 2011).
3. *The applicant withdrawal literature does not model time.* The best matched dataset in existence, roughly 25,000 applicants at a manufacturing company observed at four selection stages, was analysed with linear and logistic regression and reports no withdrawal rates, no timing data and no survival analysis (Giumetti 2011). Nobody has run a hazard model on a light industrial hiring funnel.
4. *Completed first shift does not appear as a funnel outcome.* Every source we verified terminates at offer or at hire. In temporary staffing the delivered unit is a worked shift, and the hire-to-first-shift attrition step is absent from the peer-reviewed record. The only figures available on interview no-shows and early attrition in hourly roles are vendor claims without published methodology (Workstream 2026).

The absence is not for want of data collection. The largest published recruitment-marketing benchmark dataset covers more than 1,300 employers, 379 million job-advertisement clicks and over 30 million applications, and reports an overall apply rate of 6.1 per cent at the end of 2024 against a cost per hire of \$851 (Appcast 2025). It contains no light industrial or warehouse apply rate and, more to the point here, no measure of time to first contact. The industry has instrumented the top of the funnel at enormous scale and left the interval that follows unrecorded.

A fifth absence is worth noting because it bounds what any Canadian study can say. Statistics Canada publishes the number of days a vacancy has been under active recruitment, and reports that 28.5 per cent of Canadian vacancies had been recruited against for 90 days

or more in the fourth quarter of 2025 (Statistics Canada 2026). It does not publish a fill probability conditional on that duration. There is no Canadian vacancy survival curve in the public data.

4 Candidate availability as a time-dependent quantity

4.1 Three premises

The model rests on three claims, each of which is separately supported.

Supply arrives front-loaded. Of applications to an online posting, 41 per cent arrive within the first 48 hours and 56 per cent within 96 hours (Restrepo-Echavarría and Samaniego de la Parra 2017). The pool an employer will ever see is therefore assembled early. This has a consequence that is easy to miss: the applicants who arrive first are the ones exposed longest to a slow process, so a queue worked in order of arrival applies the largest decay to the candidates who acted fastest.




Figure 2. Supply arrives front-loaded. Two in five applications to a posting land inside the first two days. The pool an employer will ever see is therefore assembled early, which means the earliest applicants are also the ones exposed longest to a slow process. Only the two marked points are reported values; the curve between and beyond them is interpolation.

Source: Restrepo-Echavarría and Samaniego de la Parra (2017), “How quickly do people apply for jobs?” Federal Reserve Bank of St. Louis. Central bank research publication. Nearly 80 million applications linked to almost 10 million postings, in technology, software, financial services, consulting and engineering occupations. Not light industrial: see the transfer caveat in section 10.

The applicant is not exclusively yours. On US federal survey microdata, jobseekers submit a mean of 13.67 applications, obtain a mean of 1.93 interviews, and face an overall probability of receiving any offer of 26.24 per cent, rising to about 37 per cent conditional on at least one interview (Dalton and Groen 2020). Roughly six applications are required per interview. An employer holding an application is holding one claim among fourteen.




Figure 3. The jobseeker runs a portfolio, not a single application. On US federal survey microdata the average jobseeker submits roughly fourteen applications over a search spell, converts about one in six into an interview, and ends with roughly a one-in-four chance of any offer. An employer is therefore never the only claim on an applicant’s attention, which is the structural reason latency has a price.

Source: Dalton and Groen (2020), *Beyond the Numbers* 9(14), US Bureau of Labor Statistics. National statistical agency. Reported: mean 13.67 applications, mean 1.93 interviews, a 0.168 interview-to-application ratio, a 26.24 per cent overall probability of receiving an offer, and a 36.89 per cent offer probability conditional on at least one interview. The report also states that “it took jobseekers on average six applications to obtain one interview.” The counts are per search spell, not simultaneously live applications, and the third bar is a probability plotted on the same scale for continuity.

Search intensity itself declines. Within an individual search spell, using online application data with controls for jobseeker composition and the evolution of available openings, search intensity “declines continuously” (Faberman and Kudlyak 2019). The applicant is a decaying asset from their own side too. We note the qualification immediately: in a German institutional setting with benefit exhaustion, search effort is roughly flat early, rises approaching exhaustion, then declines (DellaVigna et al. 2022). Search-effort decay is not universal, and the availability claim in this paper should be read as a claim about employer-side latency, not as a corollary of generic search decay.

4.2 Formal statement

Let an application arrive at time zero and let t denote elapsed time until the employer’s first substantive contact. Write $S(t)$ for the probability that the applicant is still available at t , meaning not yet withdrawn and not yet committed elsewhere, and π for the probability that an available applicant contacted at any time converts to a completed first shift. Expected conversions from a pool of N applications processed with latency distribution G are then:

$$(1) \quad E[\text{conversions}] = N \cdot \pi \cdot \int S(t) dG(t)$$

Two things follow. First, conversions are linear in N but non-linear in the latency distribution, so an operation that raises applications while its latency distribution shifts right can lose ground. Second, and more usefully, the derivative with respect to a uniform reduction in latency depends entirely on the shape of S . If S is convex, effort spent compressing the left

tail of G is worth far more than effort spent trimming the right tail, and the reverse holds if S is concave. This is the practical content of the convexity question, and it is why “respond faster” is not an operating instruction until the shape is known.

4.3 Three shapes, one prescription

Current guidance is compatible with at least three specifications of S , which imply materially different operating rules.

1. *Exponential decay.* $S(t) = \exp(-\lambda t)$. A constant proportional loss per unit time. Every hour saved is worth the same proportional amount, wherever it is saved. Under this shape a uniform improvement in average response time is the right target.
2. *Power-law decay.* $S(t) \propto (1 + t/a)^{-b}$. Steep early, nearly flat later. The first hour is worth many later hours, and an applicant already three days old is nearly as good, or nearly as bad, as one five days old. Under this shape the right target is the fastest decile, and improving the average by improving the slow tail is close to worthless.
3. *Threshold.* S approximately flat, then a discontinuity at some deadline. There is a race to beat and no reward for beating it by more. Under this shape the right target is the share of applications contacted inside the deadline, and average response time is the wrong metric entirely.




Figure 4. Three functional forms that the practitioner literature does not distinguish, and that a hazard study would. All three are consistent with the loose claim that “faster is better,” and each implies a different operating rule. Under exponential decay every saved hour is worth the same proportional amount. Under a power law the first hour is worth many later hours and the tail is nearly flat. Under a threshold there is a deadline to beat and no value in beating it by more. These are candidate specifications, not fitted curves, and no data underlies them.

Source: Specifications constructed for this paper. The two-day threshold is drawn as an illustrative example because it is the interval most often asserted in practitioner guidance; see section 3.1.

These are not competing theories so much as competing operating regimes, and the choice between them is empirical. The reason it matters is that the three recommend opposite allocations of the same scarce recruiter hour. No published work discriminates among them.

5 Evidence from adjacent literatures

5.1 Convex decay against response latency

The strongest quantitative statement of front-loaded decay against response latency that we could verify anywhere comes from outside hiring. An audit of 2,241 firms measured how long each took to respond to a web-generated enquiry: 37 per cent responded within an hour, 16 per cent within one to 24 hours, 24 per cent took longer than a day, and 23 per cent never responded at all. Among those that did respond within 30 days, the average was 42 hours. Across 1.25 million enquiries at 42 firms, those contacting within an hour were “nearly seven times as likely to qualify the lead...as those that tried to contact the customer even an hour later” and “more than 60 times as likely as companies that waited 24 hours or longer” (Oldroyd et al. 2011).




Figure 5. Employer response latency to an inbound enquiry, audited rather than self-reported. Almost half of the audited firms either took longer than a day or never replied at all. The distribution is from business enquiry handling rather than job applications, and it is the closest audited latency distribution the published record contains.

Source: Oldroyd, McElheran and Elkington (2011), “The short life of online sales leads,” *Harvard Business Review*. Practitioner magazine, not peer reviewed. Not peer reviewed. Percentages as published; they sum to 100.

We report this with two explicit qualifications. It is a practitioner magazine article rather than a peer-reviewed study, and its outcome is a qualified sales lead rather than a hire. Transposing it to hiring is precisely the untested step, and it is the step this paper proposes to test rather than assume. Its value here is that it establishes the shape is not merely conceivable: a seven-fold difference across one hour and a sixty-fold difference across a day is a convexity strong enough that, if anything like it holds in hiring, it would dominate every other lever in a recruiting operation.

5.2 Convex decay in a hiring setting

Convex decay with a flat tail has been measured in hiring, in a resume audit study that sent roughly 12,000 resumes to roughly 3,000 postings across 100 cities. About 4.7 per cent of resumes received a callback. Callbacks at eight months of unemployment were about 45 per cent lower than at one month, falling from roughly 7 per cent to roughly 4 per cent, and beyond eight months the marginal effect of additional months was “negligible” (Kroft et al. 2013). Duration dependence was stronger in tighter labour markets.




Figure 6. Convex decay with a flat tail, in the one hiring setting where the shape has been measured. Callback probability falls steeply over the first eight months of a jobseeker’s unemployment spell and then stops falling. The clock here is the applicant’s unemployment duration, not the employer’s response latency, so this is a structural analogue rather than direct evidence. The plotted curve interpolates between the two reported anchors and the authors’ description of the tail; only the marked points and the mean are reported values.

Source: Kroft, Lange and Notowidigdo (2013), *Quarterly Journal of Economics* 128(3). Peer reviewed. Resume audit study, roughly 12,000 resumes sent to roughly 3,000 postings across 100 cities. Reported: a 4.7 per cent mean callback rate, callbacks at eight months “about 45 percent lower than at one month,” and thereafter a “negligible” marginal effect of further months.

This is the clearest published example of the functional form in question, and its clock is the wrong one. It measures how an employer discounts an applicant as the applicant’s unemployment lengthens, not how an applicant discounts an employer as the employer’s silence lengthens. We present it as a structural analogue: in the one hiring context where a probability has been traced against elapsed time, the shape was steep then flat, which is the power-law regime rather than the exponential one.

5.3 Recruiting intensity as the operative margin

Labour economics contains a construct that comes close to the one this paper needs. Using establishment-level data from the Job Openings and Labor Turnover Survey, vacancy fill rates were found to rise sharply with employer growth, fall with establishment size, and rise with worker turnover, and a generalised matching function with a *recruiting intensity* margin

was shown to improve fit (Davis et al. 2013). The relevance is conceptual rather than numerical: recruiting intensity is the licence to treat something other than vacancy count as the operative variable. Response latency is an intensity margin, and, unlike most, it is one an operator can move within a week.

6 Does moving faster cost validity?

The most serious objection to compressing a hiring funnel is that speed buys volume at the cost of judgment. The evidence does not support the objection in the form it is usually made, and the reason is that the variable which carries validity is structure, not duration.

The classic meta-analysis of employment interview validity, covering 106 structured samples with 12,847 participants and 39 unstructured samples with 9,330, found structured interviews more valid than unstructured, situational interviews more valid than job-related ones, and, notably for a high-volume operation, individual interviews ($\rho = .43$) more valid than board interviews ($\rho = .32$) (McDaniel et al. 1994).

Those estimates have since been revised downward. A re-analysis of the corrections used across the selection literature concluded that the standard range-restriction adjustments “often result in substantial overcorrection” and that validity had been “substantially overestimated,” with mean estimates reduced by .10 to .20 points (Sackett et al. 2022). On the revised figures, structured interviews emerged as the highest-validity widely used predictor at $\rho = .42$, against .19 for unstructured interviews (Sackett et al. 2023).




Figure 7. Structure, not duration, is what carries interview validity. On the current best estimates, a structured interview predicts job performance more than twice as well as an unstructured one, and better than any other widely used predictor. Nothing in this ranking rewards taking longer. The implication for a high-volume funnel is that speed and rigour are not opposed: the design choice that raises validity is standardisation, and standardisation is also what makes a screen fast.

Source: Sackett, Zhang, Berry and Lievens (2023), *Industrial and Organizational Psychology* 16, table 1. Peer reviewed. These estimates are deliberately lower than older summaries: the authors find prior meta-analyses corrected systematically for range restriction and that validity had been “substantially overestimated,” with means reduced by .10 to .20.

Read operationally, that gap of .23 is the largest single design lever available in a screening process, and it is available only through standardisation. Standardisation is also what makes a screen fast, because a fixed question set with fixed scoring can be run at a fixed cost per candidate whereas a conversation cannot. Speed and validity are complements here, not substitutes. Evidence from job testing across 15 employers in the low-skilled service sector points the same way: managers who overrode test recommendations produced worse outcomes on average, which implies the overrides reflected error rather than superior private information (Hoffman et al. 2018).

One genuine caution remains, and it concerns mode rather than speed. Two meta-analyses of technology-mediated interviewing disagree. The earlier, across 13 samples and 1,557 participants, found interview ratings and applicant reactions both lower in mediated interviews ($d = -.41$ and $-.36$ respectively) (Blacksmith et al. 2016). The later, across 31 samples from 24 studies, finds essentially no effect on interview ratings ($\rho = -.01$), with the cost appearing instead as reduced organisational attractiveness ($\rho = -.17$) and reduced behavioural intentions ($\rho = -.11$), alongside a positive effect on functionality-related attitudes ($\rho = +.24$) (Moon et al. 2026).




Figure 8. Two meta-analyses of the same question reach opposite conclusions about whether mediated interviewing costs measurement quality. The earlier synthesis found a substantial penalty to interview ratings; the later and larger one finds essentially none, and locates the residual cost in how attractive the employer looks rather than in how well the interview measures. Neither disaggregates telephone from video or from asynchronous formats, so no mode-specific estimate exists.

Source: Upper group: Blacksmith, Willford and Behrend (2016), *Personnel Assessment and Decisions* 2(1), reporting Cohen's d from 13 samples, $N = 1,557$. Lower group: Moon, Park and Park (2026), *Current Psychology* 45, reporting ρ from 31 samples across 24 studies. Both peer reviewed. The two groups use different effect metrics and are plotted on a shared axis for direction and rough magnitude only.

A gap that matters for any claim about telephone screening. Neither meta-analysis disaggregates telephone from videoconference from asynchronous video, and the 1994 interview meta-analysis contains no telephone moderator. There is therefore no published mode-specific

validity estimate for telephone or automated voice screening. Any claim about the validity of a phone screen is at present an extrapolation from pooled mediated-interview effects, and the two available pooled estimates disagree about whether mode costs measurement quality at all.

7 The Canadian light industrial setting

The Canadian market in which this question has practical force is small enough to describe precisely. Employment services generated \$24.7 billion in operating revenue in 2023, up 5.2 per cent, with temporary staffing services at 48.7 per cent of total sales and Ontario accounting for 53.7 per cent of revenues (Statistics Canada 2025). Job vacancies stood at 495,100 in the fourth quarter of 2025 at a vacancy rate of 2.8 per cent, with an average offered hourly wage of \$29.25 for vacant positions and an unemployment-to-vacancy ratio of 3.1 (Statistics Canada 2026).

Table 3. The Canadian setting, from official statistics. Figures are the most recent published at the time of writing.

Indicator	Value	Period	Source
<i>Labour demand</i>			
Job vacancies, all industries	495,100	Q4 2025	StatCan
Job vacancy rate	2.8%	Q4 2025	StatCan
Unemployment-to-vacancy ratio	3.1	Q4 2025	StatCan
Average offered hourly wage, vacant positions	\$29.25	Q4 2025	StatCan
Vacancies under recruitment 90 days or more	28.5%	Q4 2025	StatCan
Vacancies, administrative and support services	23,200	Oct 2025	StatCan
<i>Industry structure</i>			
Employment services operating revenue	\$24.7B	2023	StatCan
Temporary staffing share of employment services sales	48.7%	2023	StatCan
Ontario share of employment services revenue	53.7%	2023	StatCan
<i>Workforce</i>			
Employees in a temporary job	13.3%	2018	StatCan
Average hourly earnings, temporary employees	\$21.80	2018	StatCan

Notes: The temporary employment figures are from a 2018 reference year and are the most recent for which Statistics Canada has published this breakdown in an analytical release. Neither the Labour Force Survey nor the Job Vacancy and Wage Survey publishes a separate series for temporary help agency employment, so no official Canadian headcount for agency-assigned workers exists. See section 10.

Two features of the Canadian measurement apparatus bear directly on this paper. The first is that Statistics Canada’s vacancy concept is itself a recruitment-duration concept: a position counts as vacant only if “the employer is actively recruiting outside the organization to fill the job,” and vacancy duration is defined as “the number of days of active recruitment” (Stat-

istics Canada 2024). The national statistic is measured on the same clock this paper argues is mismanaged, which makes the absence of a conditional fill rate against it more striking rather than less.

The second is the regulatory definition of the delivered unit. Since 1 July 2024, temporary help agencies in Ontario must hold a licence to operate, and clients are prohibited from knowingly engaging an unlicensed agency (Ontario, Ministry of Labour n.d.). The worker is an assignment employee, and the employment relationship with the agency persists between assignments. In this structure a hire without a worked shift is not a partial delivery, it is no delivery, which is why the terminal event in the research question is a completed first shift.

8 A measurement protocol

Two definitional choices are inherited rather than invented. The terminal event is a completed first shift, which matches the definition of a filled shift used in the Canadian light industrial benchmark study now in the field, namely a shift “covered by a confirmed worker who attended” (Kordis Research n.d.). And the elapsed-time measure is defined from the applicant’s action rather than from the operator’s, in the same way that the benchmark study defines time to fill from the moment a shift becomes open rather than from the moment a coordinator notices it. Adopting those definitions rather than coining new ones is what makes the protocol below and the field measurement commensurable.

The study that would answer this question is not expensive and does not require a randomised trial, because the necessary variation is already generated by ordinary operations. What it requires is that the clock be recorded, which at present it generally is not.

Table 4. Proposed measurement protocol. The minimum event schema needed to estimate the latency-to-conversion hazard, and the identification problem each field addresses.

Field	Definition	Why it is required
<code>t_apply</code>	Timestamp at which the application record is created.	Origin of the clock. Must be the applicant’s submission, not the recruiter’s first view of it.
<code>t_contact</code>	Timestamp of the first substantive contact attempt reaching the applicant, defined as a completed voice call, a delivered message requesting a specific action, or a scheduled screen.	The treatment variable. Delivery, not dispatch: a queued message that never arrives is not contact.
<code>t_screen, t_offer, t_shift1</code>	Completion of structured screen, offer extended, first shift worked to completion.	Separates the funnel into stages so the decay can be attributed to a step rather than to the whole process.
<code>channel, attempts</code>	Modality of each attempt and the count preceding first contact.	Attempts consume wall-clock time, so effort and latency are not independent budgets (Olson 2006).
<code>queue_depth</code>	Count of unactioned applications for the same requisition at <code>t_apply</code> .	The instrument. Queue depth shifts latency without plausibly acting on the individual applicant’s intent to convert.

Field	Definition	Why it is required
req_id, site, shift_pattern, wage	Requisition, client site, shift pattern and posted wage.	Fixed effects. Slow requisitions may also be unattractive requisitions; without them the estimate is confounded.
withdrew, reason	Applicant-initiated exit and its stated reason where captured.	Distinguishes applicant withdrawal from employer rejection. The existing withdrawal literature conflates timing with neither (Giumetti 2011).

Notes: The estimand is the cause-specific hazard of applicant withdrawal as a function of elapsed latency, with conversion to a completed first shift as the competing terminal event. A discrete-time complementary log-log specification with requisition fixed effects and hour-level risk intervals is sufficient and avoids imposing a parametric shape on the decay, which is the quantity of interest.

8.1 Identification

The obvious threat is that slow requisitions are also unattractive requisitions. If a client site pays poorly, runs an awkward shift pattern and is slow to be worked by recruiters, a naive regression of conversion on latency will attribute the site’s unattractiveness to its slowness. Requisition fixed effects absorb this, at the cost of identifying the effect only from within-requisition variation in latency, which is precisely the variation an operator controls.

The instrument we propose is queue depth at the moment of application: the number of unactioned applications for the same requisition when a given application arrives. Queue depth shifts the latency an applicant experiences for reasons that are plausibly unrelated to that applicant’s own propensity to convert, because it is determined by the arrival pattern of other people. The exclusion restriction is not innocuous and should be tested rather than asserted: if a deep queue reflects a burst of interest driven by an attractive posting, it will be correlated with applicant quality.

Two further design notes. Attempt count and latency are not independent budgets, because attempts consume wall-clock time. In a large survey methodology study the mean number of calls to first contact among contacted cases was 4.29, with attempts ranging up to 102 (Olson 2006); any specification that treats effort and speed as separate covariates without modelling their dependence will mis-attribute the effect. And the estimand should be a cause-specific hazard with competing risks, because applicant withdrawal and employer rejection are different events with different policy implications, and the existing withdrawal literature does not separate them by time (Giumetti 2011).

9 Discussion

If the availability-decline model is right, three things follow for how a high-volume hourly operation should be designed, and each of them cuts against current practice.

The pool metric is the wrong metric. Applications per opening measures the stock. What determines output under a decaying S is the joint distribution of pool age at contact, which is a property of the process rather than of the market. An operation can raise applications per

opening and lose hires, as the aggregate manufacturing data appears to show (iCIMS 2025). The corresponding metric is the share of applications reaching first substantive contact within a stated interval, which is a service level rather than a volume.

Average response time may be the wrong service level. If S follows a power law, the average is dominated by the slow tail while the conversions are determined by the fast head, and improving the average by rescuing week-old applications will produce no measurable gain. This is a testable difference between the three specifications in section 4.3 and it is the practical reason the functional form matters.

Standardisation is the mechanism, not a compromise for it. The validity evidence gives no support to the view that a faster screen is a laxer one. Structured screening outperforms unstructured screening by more than a factor of two on operational validity (Sackett et al. 2023), and structure is the property that makes throughput possible. The design that is fast is also the design that predicts performance best.

There is also a labour-market reading. Roughly three in ten Canadian vacancies have been under active recruitment for 90 days or more (Statistics Canada 2026). If a material share of that duration reflects process latency rather than genuine scarcity of willing workers, then some of what is currently discussed as a labour shortage is a queueing problem, and it would be visible as one in data that nobody currently collects.

10 Limitations

The central limitation is the one this paper is about. No dataset yet exists in which application-to-contact latency and conversion to a completed first shift are both observed, so every quantitative anchor used here is transferred from another domain, and each transfer has a direction of likely bias that should be stated.

- The convex latency decay in Oldroyd et al. (2011) comes from business enquiry handling. Job applicants have fewer substitutes than business buyers in some markets and more in others, so the sign of the transfer error is not obvious. It is also not peer reviewed.
- The convex decay with a flat tail in Kroft et al. (2013) runs on unemployment duration, not employer latency. It establishes that the shape occurs in hiring; it does not establish that it occurs on this clock.
- The front-loaded arrival pattern in Restrepo-Echavarría and Samaniego de la Parra (2017) is from technology, financial services and engineering postings. Light industrial application arrival may be more or less concentrated, and no published Canadian equivalent exists.
- The funnel ratios in Dalton and Groen (2020) are US and economy-wide, not Canadian and not light industrial.
- The paradox in Figure 1 rests on vendor-published data from one platform. The pattern is consistent with the model but is not identified: falling hires with rising applications is also consistent with a demand-side slowdown, with compositional change in who

applies, or with measurement change in the underlying panel. It motivates the question and does not answer it.

- There is no official Canadian series for temporary help agency employment. Revenue is published and temporary employment incidence is published, but not agency headcount, so the population to which any Canadian estimate would generalise cannot currently be sized from public data.

A conceptual limitation is worth naming as well. This paper treats π , the conversion probability of an available applicant, as independent of latency. That is a simplification. If slow employers are read as poorly run employers, as Rynes et al. (1990) suggests, then latency reduces both availability and willingness, and the two channels cannot be separated without measuring stated reasons for withdrawal.

11 Conclusion

The applicant pool in high-volume hourly hiring is managed as a stock and behaves, we argue, as a quantity that declines in elapsed time. The aggregate evidence is at least consistent with that reading: applicant supply per manufacturing opening rose while hires fell and time to fill lengthened, which a static-pool model cannot generate. The mechanism is available in adjacent literatures. Supply arrives front-loaded, the applicant holds roughly fourteen competing claims, process conduct matters most early in the funnel, faster offers raise acceptance without a quality penalty, and the one hiring probability that has been traced against elapsed time decayed steeply and then flattened.

What is missing is the measurement. Nobody has recorded the interval between an application and the first substantive reply and tested it against conversion to a worked shift. The prescription to move faster is universal and its evidence base is close to empty. The study that would fix this requires two timestamps that most operations already generate and do not keep, and a hazard specification that does not impose the answer. Until it is run, the correct description of the industry's central operating belief is that it is plausible, well motivated, and unmeasured.

Disclosure

Funding and interest

This paper was prepared at Kordis Research. Kordis sells software to temporary staffing agencies, which are among the operations whose practices are examined here. That is a direct commercial interest in the conclusion that response latency matters, and readers should weigh the argument accordingly. No external funding was received, and no client data was used.

What is derived rather than reported

Figure 2 and Figure 6 interpolate between reported anchor points; only the marked points and stated means are published values. Figure 4 shows constructed specifications and contains no data. The model in section 4.2 is an accounting identity and is not fitted to anything. All other numeric values in figures and tables are as published by the cited source.

Typographic note

Set in STIX Two Text. There are no em dashes anywhere in the series. En dashes appear only in numeric ranges, page ranges, and quoted source titles.

References

- Appcast. 2025. *2025 recruitment marketing benchmark report* [press release]. PR Newswire. <https://www.prnewswire.com/news-releases/appcast-releases-2025-recruitment-marketing-benchmark-report-302372820.html>.
- Barattucci, M., A. Russo, J. Grobelny, G. Santisi, and T. Ramaci. 2025. “Candidates’ reactions to job application rejections at different phases of the recruitment process.” *Journal of Management & Organization* 31 (5): 2341–2359. https://ideas.repec.org/a/cup/jomorg/v31y2025i5p2341-2359_5.html.
- Bauer, T. N., D. M. Truxillo, R. J. Sanchez, J. M. Craig, P. Ferrara, and M. A. Campion. 2001. “Applicant reactions to selection: Development of the Selection Procedural Justice Scale.” *Personnel Psychology* 54 (2): 387–419. <https://doi.org/10.1111/j.1744-6570.2001.tb00097.x>.
- Becker, W. J., T. Connolly, and J. E. Slaughter. 2010. “The effect of job offer timing on offer acceptance, performance, and turnover.” *Personnel Psychology* 63 (1): 223–241. <https://doi.org/10.1111/j.1744-6570.2009.01167.x>.
- Blacksmith, N., J. C. Willford, and T. S. Behrend. 2016. “Technology in the employment interview: A meta-analysis and future research agenda.” *Personnel Assessment and Decisions* 2 (1): Article 2. <https://scholarworks.bgsu.edu/pad/vol2/iss1/2/>.
- CareerBuilder. n.d. *66% of candidates wait less than 2 weeks before moving on*. Reporting the CareerBuilder Candidate Behavior study. <https://resources.careerbuilder.com/news-research/66-of-candidates-wait-less-than-2-weeks-before-moving-on>.
- Cotton, T. 2025. *November 2025 workforce report*. iCIMS, November 26. <https://www.icims.com/blog/workforce-report-november-2025/>.
- Dalton, M. R., and J. A. Groen. 2020. “How do jobseekers search for jobs? New data on applications, interviews, and job offers.” *Beyond the Numbers* 9 (14). US Bureau of Labor Statistics. <https://www.bls.gov/opub/btn/volume-9/how-do-jobseekers-search-for-jobs.htm>.
- Davis, S. J., R. J. Faberman, and J. C. Haltiwanger. 2013. “The establishment-level behavior of vacancies and hiring.” *Quarterly Journal of Economics* 128 (2): 581–622. <https://doi.org/10.1093/qje/qjt002>.
- DellaVigna, S., J. Heining, J. F. Schmieder, and S. Trenkle. 2022. “Evidence on job search models from a survey of unemployed workers in Germany.” *Quarterly Journal of Economics* 137 (2): 1181–1232. <https://doi.org/10.1093/qje/qjab039>.
- Faberman, R. J., and M. Kudlyak. 2019. “The intensity of job search and search duration.” *American Economic Journal: Macroeconomics* 11 (3): 327–357. <https://doi.org/10.1257/mac.20170315>.

- Gilliland, S. W. 1993. "The perceived fairness of selection systems: An organizational justice perspective." *Academy of Management Review* 18 (4): 694–734. <https://doi.org/10.5465/amr.1993.9402210155>.
- Giumetti, G. W. 2011. *Applicant self-selection during the hiring process: Developing and testing a model of applicant withdrawal* [Doctoral dissertation]. Clemson University, All Dissertations 734. https://open.clemson.edu/all_dissertations/734/.
- Government of Ontario, Ministry of Labour, Immigration, Training and Skills Development. n.d. *Temporary help agencies*. Your guide to the Employment Standards Act. <https://www.ontario.ca/document/your-guide-employment-standards-act-0/temporary-help-agencies>.
- Hausknecht, J. P., D. V. Day, and S. C. Thomas. 2004. "Applicant reactions to selection procedures: An updated model and meta-analysis." *Personnel Psychology* 57 (3): 639–683. <https://doi.org/10.1111/j.1744-6570.2004.00003.x>.
- Hoffman, M., L. B. Kahn, and D. Li. 2018. "Discretion in hiring." *Quarterly Journal of Economics* 133 (2): 765–800. <https://doi.org/10.1093/qje/qjx042>.
- Kordis Research. n.d. *Canadian temp staffing operations benchmarks (2026)*. Study protocol and pre-registered definitions. <https://kordisai.com/research/canadian-light-industrial-staffing-benchmarks-2026>.
- Kroft, K., F. Lange, and M. J. Notowidigdo. 2013. "Duration dependence and labor market conditions: Evidence from a field experiment." *Quarterly Journal of Economics* 128 (3): 1123–1167. <https://doi.org/10.1093/qje/qjt015>.
- Long, R. 2014. *The two-day rule of recruiting*. Workable, January 8; updated 2023. <https://resources.workable.com/stories-and-insights/two-day-rule-of-recruiting>.
- McDaniel, M. A., D. L. Whetzel, F. L. Schmidt, and S. D. Maurer. 1994. "The validity of employment interviews: A comprehensive review and meta-analysis." *Journal of Applied Psychology* 79 (4): 599–616. <https://home.ubalt.edu/tmitch/645/articles/McDanieletal1994CriterionValidityInterviewsMeta.pdf>.
- Moon, R. H., J. Park, and H. I. Park. 2026. "Technology-mediated interviews: Meta-analyses on interview ratings and job applicants' responses." *Current Psychology* 45: Article 349. <https://doi.org/10.1007/s12144-025-08946-x>.
- Oldroyd, J. B., K. McElheran, and D. Elkinington. 2011. "The short life of online sales leads." *Harvard Business Review*, March. <https://hbr.org/2011/03/the-short-life-of-online-sales-leads>.
- Olson, K. 2006. "Survey participation, nonresponse bias, measurement error bias, and total bias." *Public Opinion Quarterly* 70 (5): 737–758. <https://doi.org/10.1093/poq/nfl038>.
- Restrepo-Echavarría, P., and B. Samaniego de la Parra. 2017. *How quickly do people apply for jobs? On the Economy*. Federal Reserve Bank of St. Louis, November 13. <https://www.stlouisfed.org/on-the-economy/2017/november/how-quickly-people-apply-jobs>.
- Rynes, S. L., R. D. Bretz, and B. Gerhart. 1990. *The importance of recruitment in job choice: A different way of looking*. CAHRS Working Paper 90-24. Cornell University. <http://digitalcommons.ilr.cornell.edu/cahrswp/389/>.
- Sackett, P. R., C. Zhang, C. M. Berry, and F. Lievens. 2022. "Revisiting meta-analytic estimates of validity in personnel selection: Addressing systematic overcorrection for restriction of range." *Journal of Applied Psychology* 107 (11): 2040–2068. <https://doi.org/10.1037/apl0000994>.
- Sackett, P. R., C. Zhang, C. M. Berry, and F. Lievens. 2023. "Revisiting the design of selection systems in light of new findings regarding the validity of widely used predictors." *Industrial and Organizational Psychology* 16: 283–300. <https://doi.org/10.1017/iop.2023.24>.

- Society for Industrial and Organizational Psychology. 2012. *What we know about applicant reactions to selection: Research summary and best practices* (T. N. Bauer, J. McCarthy, N. Anderson, D. M. Truxillo, and J. F. Salgado). SIOP White Paper Series. https://www.siop.org/wp-content/uploads/2024/07/SIOP-Aplicant_Reactions_to_Selection_final.pdf.
- Statistics Canada. 2019. *Temporary employment in Canada, 2018*. Catalogue 11-627-M, May 14. <https://www150.statcan.gc.ca/n1/pub/11-627-m/11-627-m2019034-eng.htm>.
- Statistics Canada. 2024. *Guide to the Job Vacancy and Wage Survey, 2024*. Catalogue 75-514-G. <https://www150.statcan.gc.ca/n1/pub/75-514-g/75-514-g2024001-eng.htm>.
- Statistics Canada. 2025. *Employment services, 2023*. The Daily, March 17. Source table 21-10-0063-01. <https://www150.statcan.gc.ca/n1/daily-quotidien/250317/dq250317d-eng.htm>.
- Statistics Canada. 2025. *Payroll employment, earnings and hours, and job vacancies, October 2025*. The Daily, December 18. <https://www150.statcan.gc.ca/n1/daily-quotidien/251218/dq251218a-eng.htm>.
- Statistics Canada. 2026. *Job vacancies, fourth quarter 2025*. The Daily, March 17. Source table 14-10-0398-01. <https://www150.statcan.gc.ca/n1/daily-quotidien/260317/dq260317a-eng.htm>.
- Uggerslev, K. L., N. E. Fassina, and D. Kraichy. 2012. "Recruiting through the stages: A meta-analytic test of predictors of applicant attraction at different stages of the recruiting process." *Personnel Psychology* 65 (3): 597–660. <https://doi.org/10.1111/j.1744-6570.2012.01254.x>.
- Workstream. 2026. *Interview no-shows are predictable. That means they're preventable*. March 19. <https://www.workstream.us/blog/predicting-and-preventing-interview-no-shows>.

Appendix A. Source classification

Every source cited in this paper, with the class assigned to it on retrieval. Full details are in the reference list.

Table A1. Source classification.

Source	Class
Appcast 2025	Vendor or practitioner publication
Barattucci, M., A. Russo, J. Grobelny, G. Santisi, and T. Ramaci 2025	Peer-reviewed
Bauer, T. N., D. M. Truxillo, R. J. Sanchez, J. M. Craig, P. Ferrara, and M. A. Campion 2001	Peer-reviewed
Becker, W. J., T. Connolly, and J. E. Slaughter 2010	Peer-reviewed
Blacksmith, N., J. C. Willford, and T. S. Behrend 2016	Peer-reviewed
CareerBuilder n.d.	Vendor or practitioner publication
Cotton, T. 2025	Vendor or practitioner publication
Dalton, M. R., and J. A. Groen 2020	Government department or agency
Davis, S. J., R. J. Faberman, and J. C. Haltiwanger 2013	Peer-reviewed
DellaVigna, S., J. Heining, J. F. Schmieder, and S. Trenkle 2022	Peer-reviewed
Faberman, R. J., and M. Kudlyak 2019	Peer-reviewed
Gilliland, S. W. 1993	Peer-reviewed
Giumetti, G. W. 2011	Doctoral dissertation

Source	Class
Government of Ontario, Ministry of Labour, Immigration, Training and Skills Development n.d.	Regulator or statute
Hausknecht, J. P., D. V. Day, and S. C. Thomas 2004	Peer-reviewed
Hoffman, M., L. B. Kahn, and D. Li 2018	Peer-reviewed
Kordis Research n.d.	Author's own research programme
Kroft, K., F. Lange, and M. J. Notowidigdo 2013	Peer-reviewed
Long, R. 2014	Vendor or practitioner publication
McDaniel, M. A., D. L. Whetzel, F. L. Schmidt, and S. D. Maurer 1994	Peer-reviewed
Moon, R. H., J. Park, and H. I. Park 2026	Peer-reviewed
Oldroyd, J. B., K. McElheran, and D. Elkington 2011	Vendor or practitioner publication
Olson, K. 2006	Peer-reviewed
Restrepo-Echavarría, P., and B. Samaniego de la Parra 2017	Government department or agency
Rynes, S. L., R. D. Bretz, and B. Gerhart 1990	Working paper
Sackett, P. R., C. Zhang, C. M. Berry, and F. Lievens 2022	Peer-reviewed
Sackett, P. R., C. Zhang, C. M. Berry, and F. Lievens 2023	Peer-reviewed
Society for Industrial and Organizational Psychology 2012	Industry body
Statistics Canada 2019	National statistical agency
Statistics Canada 2024	National statistical agency
Statistics Canada 2025	National statistical agency
Statistics Canada 2025	National statistical agency
Statistics Canada 2026	National statistical agency
Uggerslev, K. L., N. E. Fassina, and D. Kraichy 2012	Peer-reviewed
Workstream 2026	Vendor or practitioner publication

Notes: Assigned on retrieval and not revised. Where a source appears in more than one paper in this series it carries the same class in each.

Kordis Research Working Paper Series · KRS-2026-01 · Version 1.0, August 2026 · Patrick Underwood

Not peer reviewed. Cite as: Underwood, P. (2026). *Response Latency and Candidate Availability in High-Volume Hourly Hiring* (Kordis Research Working Paper KRS-2026-01).